

EXHIBIT A

Data ONTAP® 7.2

Network Management Guide

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>

Part number 210-03535_A0
Updated for Data ONTAP 7.2.2 on 22 March 2007

Copyright and trademark information

Copyright information

Copyright © 1994–2007 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Portions of this product are derived from the Berkeley Net2 release and the 4.4-Lite-2 release, which are copyrighted and publicly distributed by The Regents of the University of California.

Copyright © 1980–1995 The Regents of the University of California. All rights reserved.

Portions of this product are derived from NetBSD, copyright © Carnegie Mellon University.

Copyright © 1994, 1995 Carnegie Mellon University. All rights reserved. Author Chris G. Demetriou.

Permission to use, copy, modify, and distribute this software and its documentation is hereby granted, provided that both the copyright notice and its permission notice appear in all copies of the software, derivative works or modified versions, and any portions thereof, and that both notices appear in supporting documentation.

CARNEGIE MELLON ALLOWS FREE USE OF THIS SOFTWARE IN ITS “AS IS” CONDITION. CARNEGIE MELLON DISCLAIMS ANY LIABILITY OF ANY KIND FOR ANY DAMAGES WHATSOEVER RESULTING FROM THE USE OF THIS SOFTWARE.

Software derived from copyrighted material of The Regents of the University of California and Carnegie Mellon University is subject to the following license and disclaimer:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notices, this list of conditions, and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display this text:
This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER

IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software contains materials from third parties licensed to Network Appliance Inc. which is sublicensed, and not sold, and title to such material is not passed to the end user. All rights reserved by the licensors. You shall not sublicense or permit timesharing, rental, facility management or service bureau usage of the Software.

Portions developed by the Apache Software Foundation (<http://www.apache.org/>). Copyright © 1999 The Apache Software Foundation.

Portions Copyright © 1995–1998, Jean-loup Gailly and Mark Adler

Portions Copyright © 2001, Sitaka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

Portions Copyright © 1995 University of Southern California. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by the University of Southern California, Information Sciences Institute. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

Portions of this product are derived from version 2.4.11 of the libxml2 library, which is copyrighted by the World Wide Web Consortium.

Network Appliance modified the libxml2 software on December 6, 2001, to enable it to compile cleanly on Windows, Solaris, and Linux. The changes have been sent to the maintainers of libxml2. The unmodified libxml2 software can be downloaded from <http://www.xmlsoft.org/>.

Copyright © 1994–2002 World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. <http://www.w3.org/Consortium/Legal/>

Software derived from copyrighted material of the World Wide Web Consortium is subject to the following license and disclaimer:

Permission to use, copy, modify, and distribute this software and its documentation, with or without modification, for any purpose and without fee or royalty is hereby granted, provided that you include the following on ALL copies of the software and documentation or portions thereof, including modifications, that you make:

The full text of this NOTICE in a location viewable to users of the redistributed or derivative work.

Any pre-existing intellectual property disclaimers, notices, or terms and conditions. If none exist, a short notice of the following form (hypertext is preferred, text is permitted) should be used within the body of any redistributed or derivative code: "Copyright © [date-of-software] World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. <http://www.w3.org/Consortium/Legal/>"

Notice of any changes or modifications to the W3C files, including the date changes were made.

THIS SOFTWARE AND DOCUMENTATION IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

COPYRIGHT HOLDERS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE SOFTWARE OR DOCUMENTATION.

The name and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to the software without specific, written prior permission. Title to copyright in this software and any associated documentation will at all times remain with copyright holders.

Software derived from copyrighted material of Network Appliance, Inc. is subject to the following license and disclaimer:

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Network Appliance.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, DataFabric, Data ONTAP, FAServer, FilerView, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, SyncMirror, Topio, VFM, and WAFL are registered trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries. Cryptainer, Cryptoshred, Datafort, and Decru are registered trademarks, and Lifetime Key Management and OpenKey are trademarks, of Decru, a Network Appliance, Inc. company, in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The evolution of storage are trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. ApplianceWatch, BareMetal, Camera-to-Viewer, ComplianceClock, ComplianceJournal, ContentDirector, ContentFabric, EdgeFiler, FlexClone, FlexShare, FPolicy, HyperSAN, InfoFabric, LockVault, NOW, NOW NetApp on the Web, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simplicore, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, SohoFiler, SpinMirror, SpinRestore, SpinShot, SpinStor, StoreVault, vFiler, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the United States and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the U.S.A.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	ix
Chapter 1	Network Interface Configuration	1
	Understanding the network interfaces on your storage system.	2
	Understanding frame size, MTU size, and jumbo frames	5
	Understanding Ethernet media types	8
	Understanding flow control.	10
	Configuring network interfaces.	12
	Configuring 10 gigabit Ethernet TOE cards	18
	Configuring aliases for an interface	24
	Changing the status of an interface to Up or Down	26
	Displaying network interface information	27
	Diagnosing network problems	29
Chapter 2	ATM Configuration	31
	About ATM and ATM LANE	32
	Preparing the ATM adapter for LANE	37
	Verifying that the ATM adapter is installed and functioning	39
	Verifying a working connection to the ATM network	40
	Verifying that the UNI is operational	41
	Configuring the LANE Configuration Server address	43
	Configuring the ATM adapter for an Emulated LAN	46
	Adding an Emulated LAN to the ATM adapter	47
	Configuring the logical Ethernet interface	49
	Deleting an Emulated LAN from an ATM adapter	50
	Checking and completing the Emulated LAN configuration.	51
	Verifying the communications link	52
	Checking the configuration settings	53
	Checking the other elements of the Emulated LAN	54
	Modifying load balancing and failover	56
	Saving the ATM configuration commands in the /etc/rc file	58
	Saving the host and IP address data in the /etc/hosts file	59
	Understanding FORE/IP over SPANS	60

	Managing FORE/IP and PVCs	61
	Establishing FORE/IP PVCs on your storage system	62
	Displaying information about a FORE/IP PVC	64
	Displaying the FORE/IP configuration	65
	Changing the ATM adaptation layer for FORE/IP and SPANS	67
	Deleting a FORE/IP PVC	68
Chapter 3	Network Routing Configuration	69
	About routing in Data ONTAP	70
	About fast path	71
	About the routing table.	73
	Enabling and disabling routing mechanisms	76
	Displaying the routing table and default route information	78
	Modifying the routing table.	81
	Protecting your storage system from forged ICMP redirect attacks	82
	Diagnosing ping problems	83
Chapter 4	Host-Name Resolution	85
	Maintenance of host information	86
	Using the /etc/hosts file to maintain host information	87
	Using DNS to maintain host information.	91
	Using dynamic DNS to update host information	98
	Using NIS to maintain host information	101
	Changing the host name search order	110
Chapter 5	Storage System Monitoring Using SNMP	113
	Understanding SNMP implementation in Data ONTAP	114
	Understanding traps in Data ONTAP	116
	Contents of the custom MIB	119
	Contents of the iSCSI MIB.	122
	Managing the SNMP agent.	123
	Creating SNMP traps	129
	Understanding user-defined traps	130
	Defining or modifying a trap.	131
	SNMP trap parameters	136

Chapter 6	Virtual LAN (VLAN) Configuration	143
	Understanding VLANs	144
	VLANs in Data ONTAP	148
	Managing VLANs on your storage system	150
	Creating and configuring a VLAN on your storage system	151
	Adding an interface to a VLAN	154
	Deleting a VLAN.	155
	Modifying VLAN interfaces	157
	Viewing VLAN statistics.	158
 Chapter 7	 Configuring vifs	 161
	Understanding vifs	162
	Types of vifs	164
	Managing vifs.	168
	Creating a single-mode vif	170
	Selecting an active interface in a single-mode vif	172
	Creating a static or dynamic multimode vif	174
	Adding interfaces to a vif	177
	Deleting an interface from a vif	178
	Displaying the status of a vif.	179
	Displaying statistics of a vif	183
	Viewing the LACP log file.	184
	Destroying a vif	185
	Second-level vifs	186
	Understanding second-level vifs on a single storage system	187
	Creating a second-level vif on a single storage system	188
	Understanding second-level vifs in a cluster.	190
	Creating a second-level vif in a cluster	192
 Chapter 8	 Internet Protocol Security Configuration	 197
	Understanding IPsec	198
	Setting up IPsec.	203
	Managing security policies	216
	Viewing security associations	222
 Appendix A	 Network Interface Statistics	 223

	Statistics for Fast Ethernet interfaces224
	Statistics for Gigabit Ethernet and Ethernet Controller IV interfaces228
	Statistics for 10 Gigabit Ethernet interface233
	Statistics for FAS250/ FAS270 network interfaces236
	Statistics for FAS3000/V3000 series or FAS6000/V6000 series interfaces .	.240
	Statistics for ATM interfaces244
Appendix B	Improving storage system performance245
Appendix C	IP port usage on a storage system247
Appendix D	Netdiag Error Codes261
	Index267

Host-Name Resolution

4

About this chapter This chapter discusses how you can use the Data ONTAP configuration files, Domain Name System (DNS), and Network Information Service (NIS) to resolve host names.

Topics in this chapter This chapter discusses the following topics:

- ◆ “Maintenance of host information” on page 86
- ◆ “Using the /etc/hosts file to maintain host information” on page 87
- ◆ “Using DNS to maintain host information” on page 91
- ◆ “Using dynamic DNS to update host information” on page 98
- ◆ “Using NIS to maintain host information” on page 101
- ◆ “Changing the host name search order” on page 110

Maintenance of host information

Ways to maintain host information

Host information can be maintained in one or all of the following ways in Data ONTAP:

- ◆ In the `/etc/hosts` file on your storage system's default volume
For detailed information, see "Using the `/etc/hosts` file to maintain host information" on page 87.
- ◆ On a Domain Name System (DNS) server
For detailed information, "Using DNS to maintain host information" on page 91.
- ◆ On a Network Information Service (NIS) server
For detailed information, see "Using NIS to maintain host information" on page 101.

Search order for host information

If you use more than one of the above ways to maintain host information, the ways are used in the order determined by the `/etc/nsswitch.conf` file. For detailed information about this file, see "Changing the host name search order" on page 110.

The role of host-name resolution in Data ONTAP

Data ONTAP relies on correct host-name resolution to provide basic connectivity for storage systems on the network, including

- ◆ Processing NFS mount requests
- ◆ Establishing CIFS sessions
- ◆ Authenticating Remote Shell (RSH) protocol sessions to storage systems

If you are unable to access storage system data or establish sessions, there might be problems with host-name resolution on your storage system or on a name server.

Using the /etc/hosts file to maintain host information

About the /etc/hosts file

Data ONTAP uses the /etc/hosts file to resolve host names to IP addresses, including host names used in any of the following files:

- ◆ /etc/rc
- ◆ /etc/syslog.conf
- ◆ /etc/exports
- ◆ /etc/netgroup
- ◆ /etc/hosts.equiv

You must ensure that the /etc/hosts file is kept up-to-date. If you update the file, you do not need to reboot your storage system—the changes to the file take effect immediately.

When Data ONTAP is first installed, the /etc/hosts file is automatically created with default entries for the following interfaces:

- ◆ localhost
- ◆ All interfaces on your storage system

Note

The /etc/hosts file resolves the host names for the storage system it is configured on. This file cannot be used by other systems for name resolution.

For more information on file format, see the `na_hosts(5)` man page.

Ways to add entries to the /etc/hosts file

You can add IP address and hostname entries in the /etc/hosts file in the following two ways:

- ◆ Locally

You might want to add entries to the local /etc/hosts file if the number of entries is small. You can do so in the following ways:

- ◆ At the command line
See “Editing the /etc/hosts file manually” on page 88.
- ◆ Using FilerView
See “Editing the /etc/hosts file with FilerView” on page 89.

◆ Remotely using the NIS makefile master

If the number of entries is large and you have access to an NIS makefile master, you might want to use the makefile master to create the `/etc/hosts` file. This method prevents errors that could be introduced in the manual creation process. For details, see “Creating `/etc/hosts` from the NIS master” on page 89.

Note

Using NIS to distribute the `/etc/hosts` file is different from looking up host names on an NIS server. For more information about network lookups, see “Using NIS to maintain host information” on page 101.

`/etc/hosts` file hard limits

The following are hard limits for the `/etc/hosts` file:

- ◆ Maximum line size is 1022 characters.
- ◆ Maximum number of aliases is 34.
- ◆ There is no file size limit.

Note

The line size limit includes the end of line character. You can enter up to 1021 characters per line.

Editing the `/etc/hosts` file manually

To edit the `/etc/hosts` file manually, complete the following steps.

Step	Action
1	From a workstation that has access to your storage system’s root volume, open the <code>/etc/hosts</code> file using a text editor.
2	Edit the file to your needs. The format of the file is as follows: <i>IP address Host-name aliases</i>
3	Save the file.

Example: The following shows how the entries might look in the `/etc/hosts` file on a storage system:

```

192.16.3.145    toaster    toaster-e0
192.16.4.155    toaster-e2
192.16.5.165    toaster-e4
192.16.6.175    toaster-e8

```

In the first line, your storage system's host name itself is used as an alias for the first network interface. That is, network traffic addressed to toaster will be received on the toaster-e0 interface.

Editing the /etc/hosts file with FilerView

To edit the /etc/hosts file with FilerView, complete the following steps.

Step	Action
1	In FilerView, click Network in the list on the left.
2	In the list under Network, click Manage Hosts File.
3	Click in the hosts window, then click Insert.
4	Complete the fields in the Create a New /etc/hosts Line window for each host you wish to add and click OK.
5	Click Apply in the Manage Hosts File window.

Creating /etc/hosts from the NIS master

To modify the makefile for the NIS master to create a hosts file and copy it to the /etc directory on your storage system's default volume, complete the following steps.

Step	Action
1	On the NIS server, open the NIS makefile with an editor.
2	Locate the section for hosts.time.

Step	Action
3	<p>Add the following lines at the end of the hosts.time section, replacing <i>dirname</i> with a directory name of your choice, and <i>toaster 1</i>, <i>toaster2</i>, and so on with names of your storage systems:</p> <pre> @mntdir=/tmp/dirname_etc_mnt_\$\$\$;\ if [! -d \$mntdir]; then rm -f \$mntdir; \ mkdir \$mntdir; fi;\ for s_system in toaster1 toaster2 toaster3 ; do \ mount \$\$s_system:/etc \$mntdir;\ mv \$mntdir/hosts \$mntdir/hosts.bak;\ cp /etc/hosts \$mntdir/hosts;\ umount \$mntdir;\ done;\ rmdir \$mntdir </pre>
4	<p>Save the NIS makefile.</p> <p>The /etc/hosts file on your storage system is updated whenever the NIS makefile is run.</p>

/etc/netgroup file hard limits

When editing the /etc/netgroup file, please observe these hard limits:

- ◆ Maximum entry size is 4096.
- ◆ Maximum netgroup nesting limit is 1000.
- ◆ There is no file size limit.

Note

The entry size limit includes the end of line character. You can add up to 4095 characters per entry.

Using DNS to maintain host information

Advantage of using DNS

DNS enables you to maintain host information centrally. As a result, you do not have to update the `/etc/hosts` file every time you add a new host to the network. If you have several storage systems on your network, maintaining host information centrally saves you from updating the `/etc/hosts` file on each storage system every time you add or delete a host.

A conventional storage system policy for efficient host-name resolution is to do both of the following:

- ◆ Maintain a short `/etc/hosts` file containing local interfaces, as described in “Ways to add entries to the `/etc/hosts` file” on page 87.
- ◆ Enable DNS with DNS caching, as described in “About configuring DNS” on page 91 and “What DNS name caching does” on page 95.

About configuring DNS

You can configure your storage system to use one or more DNS servers either during the setup procedure or later using the command line or FilerView.

If you configure DNS during the setup procedure, your storage system’s DNS domain name and name server addresses are configured

- ◆ Automatically if you use Dynamic Host Configuration Protocol (DHCP) to configure onboard interfaces
- ◆ Manually if you do not use DHCP—you must enter the values when prompted

If you configure DNS later, you need to take these actions:

- ◆ Specify DNS name servers.
- ◆ Specify the DNS domain name of your storage system.
- ◆ Enable DNS on your storage system.

You can enable DNS and set DNS configuration values in either of these ways:

- ◆ Using FilerView
See “Configuring DNS with FilerView” on page 92.

- ◆ At the command line
 - See the appropriate instructions:
 - ◆ “Creating or editing `/etc/resolv.conf`” on page 94
 - ◆ “Specifying the DNS domain name” on page 94
 - ◆ “Disabling or enabling DNS” on page 95

If you want to use primarily DNS for host-name resolution, specify it ahead of other methods in the hosts map in the `/etc/nsswitch.conf` file. For information about how to edit the `nsswitch.conf` file, see “Changing the host name search order” on page 110.

Correct host-name resolution depends on the correct configuration of the DNS server. If you experience problems with host-name resolution or data availability, check the DNS server in addition to local networking.

For more information about storage system DNS resolution of host names, see the `na_dns(8)` man page.

Configuring DNS with FilerView

To set or modify DNS configuration values with FilerView, complete the following steps.

Step	Action
1	In FilerView, click Network in the list on the left.
2	In the list under Network, click Manage DNS and NIS Name Service.

Step	Action	
3	If you want to....	Then...
	Enable DNS	Select Yes in the DNS Enabled field.
	Set or modify the DNS domain name	Enter a name in the DNS Domain Name field. Examples of configuration values are listed in “Specifying the DNS domain name” on page 94.
	Specify or modify DNS servers	Enter up to three IP addresses in the DNS Servers fields. Examples of configuration values are listed in “Creating or editing /etc/resolv.conf” on page 94.
	Specify or modify the search list for host name lookup	Enter a name in the DNS Domain Search field.

**Creating or editing
/etc/resolv.conf**

To create or edit the /etc/resolv.conf file, complete the following step.

Step	Action	
1	If...	Then...
	You are creating the /etc/resolv.conf file	Using a text editor, create the /etc/resolv.conf file in the root volume. The file can consist of up to three lines, each specifying a name server host in the following format: nameserver <i>ip_address</i> Example: nameserver 192.9.200.10 nameserver 192.9.200.20 nameserver 192.9.200.30
	You are editing the /etc/resolv.conf file	From a workstation that has access to your storage system's root volume, edit the /etc/resolv.conf file using a text editor.

You can optionally set or modify the domain search list for DNS host name lookup. For more information, see the na_resolv.conf(5) man page

**/etc/resolv.conf
hard limits**

The following are the NFS hard limits for the /etc/resolv.conf command.

- ◆ Maximum line size is 256.
- ◆ Maximum number of name servers is 3.
- ◆ Maximum domain name length is 256.
- ◆ Maximum search domains limit is 6. The total number of characters for all search domains cannot exceed 256.
- ◆ No file size limit.

Note

The line size limit includes the end of line character. You can add up to 255 characters per line.

**Specifying the DNS
domain name**

To specify or change the DNS domain name, complete the following step at your storage system command line.

Step	Action
1	<p>Enter the following command:</p> <pre>options dns.domainname domain</pre> <p><i>domain</i> is the new domain name, which follows your storage system's host name in the fully qualified domain name.</p> <p>For example, the domain name of the storage system system1.company.com is company.com.</p>

Disabling or enabling DNS

To disable or enable DNS, complete the following step at your storage system command line.

Step	Action
1	<p>Enter the following command:</p> <pre>options dns.enable {off on}</pre> <p>Use <i>off</i> to disable DNS or <i>on</i> to enable DNS.</p>

If you did not configure DNS during the Data ONTAP setup procedure, DNS is disabled by default.

Once enabled, DNS should be disabled only when you change host-name resolution procedures or when you troubleshoot problems with the DNS name server or Windows Active Directory server.

Note

Your storage system's CIFS implementation depends on DNS to provide the Windows Active Directory service. Therefore, disabling DNS might interrupt CIFS services.

What DNS name caching does

DNS name caching enables the DNS name resolver to speed up the process by which it converts host names into IP addresses. DNS name caching stores DNS requests by caching them so that they are easy to find the next time. Name caching improves DNS performance in the case of name server failure as well as reducing the time it takes for cluster takeover and giveback.

DNS name caching is enabled by default.

Disabling or enabling DNS name caching

To disable or enable DNS name caching, complete the following step at your storage system command line.

Attention

Disabling DNS name caching clears the DNS name cache.

Step	Action
1	<p>Enter the following command:</p> <pre>options dns.cache.enable {on off}</pre> <p>Use on to enable DNS name caching or off to disable DNS name caching.</p>

Flushing the DNS cache

Entries in the DNS cache have a set expiration. If an entry that has expired is needed again, your storage system contacts the DNS server to get an updated entry. However, if a DNS entry changes before it has expired, you must flush the DNS cache to force the storage system to get the new DNS record.

If some of your DNS records change often, you should make sure that your DNS server transmits them with a low Time To Live (TTL). (You set the TTL in the DNS server.) You can also disable DNS caching on your storage system with the `dns.cache.enable` option, but doing so might reduce performance.

To flush the DNS cache, complete the following step.

Step	Action
1	<p>Enter the following command:</p> <pre>dns flush</pre>

Displaying DNS information

You can display the following types of DNS information:

- ◆ Status of the DNS resolver
- ◆ List of DNS servers configured in the `/etc/resolv.conf` file
- ◆ State of each DNS server

- ◆ Timestamp when the DNS server was last polled
- ◆ Average round-trip time of a DNS query
- ◆ Total number of DNS queries made
- ◆ Number of failed DNS queries
- ◆ Default domain configured on your storage system
- ◆ List of other domains that will be used with unqualified names for name lookup

To display DNS information, complete the following step.

Step	Action
1	Enter the following command: dns info

For more information about the `dns info` display, see the `na_dns(1)` man page.

Using dynamic DNS to update host information

About dynamic DNS updates

Dynamic DNS updates enable your storage system to send new or changed DNS information to the primary master DNS server for your storage system's zone.

Need for dynamic DNS updates

Without dynamic DNS updates, system administrators have to manually add DNS information (DNS name and IP address) to the identified DNS servers when a new system is brought online or when existing DNS information changes. This process is not only slow, but also error-prone.

Additionally, in a disaster-recovery situation when a storage system with a large number of vFiler units is brought online, manual configuration of DNS information for those vFiler units can result in a longer-than-needed downtime.

By enabling dynamic DNS updates on your storage system, you allow your storage system to automatically send information to the DNS servers as soon as the information changes on the system.

For example, if you want to change the IP address on interface e0 of StorageSystem1, you can simply configure e0 with the new IP address. StorageSystem1 automatically sends updated information to primary master DNS server for StorageSystem1.

How dynamic DNS updates work in Data ONTAP

If dynamic DNS updates are enabled on your storage system, it periodically sends updates to the primary master DNS server for its zone. Your storage system finds out the primary master DNS server for its zone by querying the DNS servers configured in storage system's `/etc/resolv.conf` file. The primary master DNS server might be different from the ones configured in your storage system's `/etc/resolv.conf` file.

By default, periodic updates are sent every 12 hours. A time-to-live (TTL) value is assigned to every DNS update sent from your storage system. The TTL value defines the time for which a DNS entry is valid on the DNS server. By default, the TTL value is set to 24 hours, and you can change it.

In addition to periodic updates, DNS updates are also sent if any DNS information changes on your storage system.

When your storage system sends an update to the DNS server, it waits up to five minutes to receive an acknowledgement of the update from the server. If it does not receive an acknowledgement, the storage system sends the update again. This time, the storage system doubles the waiting interval (to 10 minutes), before sending the update. The storage system continues to double the waiting interval with each retry until a waiting interval of 160 minutes or TTL/2, whichever is less, is reached.

Support for dynamic DNS updates in Data ONTAP

When using dynamic DNS updates in Data ONTAP, the following conditions apply:

- ◆ By default, dynamic DNS updates are disabled in Data ONTAP.
- ◆ Dynamic DNS updates are supported on UNIX and Windows systems.
- ◆ On Windows DNS servers, secure dynamic DNS updates can be used to prevent malicious updates on the DNS servers. Kerberos is used to authenticate updates.

Even if secure dynamic DNS updates are enabled, your storage system initially tries sending updates in clear text. If the DNS server is configured to accept only secure updates, the updates sent in clear text are rejected. Upon rejection, the storage system sends secure DNS updates.

- ◆ For secure dynamic DNS updates, your storage system must have CIFS running and must be using Windows Domain authentication.
- ◆ Dynamic DNS updates *can* be sent for the following:
 - ◆ Vif and VLAN interfaces
 - ◆ vFiler units
- ◆ You cannot set TTL values for individual vFiler units. All vFiler units inherit the TTL value set for vFiler0, which is the default vFiler unit and is the same as the physical storage system.
- ◆ DHCP addresses *cannot* be dynamically updated.
- ◆ In a takeover situation, the hosting storage system is responsible for sending DNS updates for IP addresses for which it is responding.

Enabling dynamic DNS updates

To enable your storage system to send dynamic DNS updates automatically, complete the following step on your storage system.

Step	Action
1	<p>Enter the following command:</p> <pre>options dns.update.enable [off on secure]</pre> <p>Off—Disable dynamic DNS updates</p> <p>On—Enable dynamic DNS updates</p> <p>Secure—Enable secure dynamic DNS updates</p> <p>Note— Secure dynamic DNS updates are supported for Windows DNS servers only.</p>

Changing the time-to-live setting for DNS entries

To change the TTL for the DNS entries, complete the following step.

Step	Action
1	<p>Enter the following command:</p> <pre>options dns.update.ttl time</pre> <p>where <i>time</i> can be set in seconds (s), minutes (m), or hours (h) with a minimum value of 600 seconds and a maximum value of 24 hours.</p> <p>For example, to set the TTL to two hours, enter the following command:</p> <pre>options dns.update.ttl 2h</pre>

Using NIS to maintain host information

Advantage of using NIS

Like DNS, NIS enables you to centrally maintain host information. NIS provides two methods for storage system host-name resolution:

- ◆ Using a makefile master on the NIS server, which creates a `/etc/hosts` file and copies it to your storage system's default volume for local host name lookup. This method is described in "Creating `/etc/hosts` from the NIS master" on page 89.
- ◆ Using a hosts map, maintained as a database on the NIS server, which your storage system queries in a host lookup request across the network. This method is described in this section.

NIS also enables you to maintain user information. For more information, see the *Data ONTAP System Administration Guide*.

Using NIS slave for name resolution

Host-name resolution using a hosts map can have a performance impact, because each query for the hosts map is sent across the network to the NIS server. To improve performance, you can enable an NIS slave on your storage system.

The NIS slave establishes a contact with an NIS master server and does the following two tasks:

- ◆ Downloads the maps from the NIS master server. Once the maps have been downloaded, they are stored in the `/etc/yp/nis_domain_name/` directory. All NIS requests from your storage system are then serviced by the NIS slave using these maps. The NIS slave checks the NIS master every 45 minutes for any changes to the maps. If there are changes, they are downloaded.
- ◆ Listens for updates from the NIS master. When the maps on the NIS master are changed, the NIS master administrator can choose to notify all slaves. Therefore, in addition to periodically checking for updates from the NIS master, the NIS slave also listens for updates from master.

Note

The NIS slave does not respond to remote NIS client requests and thus cannot be used by other NIS clients for name lookups.

Selection of an NIS master

When the NIS slave is enabled on your storage system, the NIS servers listed with the `nis.servers` option are contacted to determine the master NIS server. The NIS master can be different from the servers listed with the `nis.servers` option. If that is the case, the servers listed with the `nis.servers` option inform the slave about the master server.

Note

Either the NIS server must have an entry in the hosts map for the master or the `/etc/hosts` file on your storage system must be able to resolve the IP address of the master. Otherwise, the NIS slave on the storage system cannot contact the master.

Guidelines for using the NIS slave

Keep the following guidelines in mind when using the NIS slave on your storage system:

- ◆ The root volume of your storage system must have sufficient space to download maps for the NIS slave. Typically, the space required in the root volume is same as the size of the maps on the NIS server.
If the root volume does not have enough space to download maps, the following occurs:
 - ◆ An error message is displayed informing you that the space on the disk is not sufficient to download or update the maps from the NIS master.
 - ◆ If the maps cannot be *downloaded*, the NIS slave is disabled. Your storage system switches to using hosts map on the NIS server for name resolution.
 - ◆ If the maps cannot be *updated*, your storage system continues to use the old maps.
- ◆ If the NIS master server was started with the `-d` option or if the `hosts.byname` and `hosts.byaddr` maps are generated with the `-b` option, your storage system must have DNS enabled, DNS servers must be configured, and the hosts entry in the `/etc/nswitch.conf` file must contain DNS as an option to use for host name lookup.

If you have your NIS server configured to do host name lookups using DNS or if you use DNS to resolve names that cannot be first resolved using the `hosts.by*` maps, using the NIS slave causes those lookups to fail, because when the NIS slave is used, all lookups are performed locally using the downloaded maps. However, if you configure DNS on your storage system as described previously, the lookups succeed.

- ◆ You can use the NIS slave for the following:
 - ❖ Vif and VLAN interfaces
 - ❖ vFiler units
 - ❖ Storage system clusters

Note

Ensure that the `nis.servers` options value is the same on both cluster nodes and that the `/etc/hosts` file on both cluster nodes can resolve the name of the NIS master server.

About configuring NIS for host lookups

You can configure your storage system to use one or more NIS servers either during the setup procedure or later using the Data ONTAP command line or FilerView.

If you configure NIS later, you need to do all of the following:

- ◆ Specify the NIS server to which your storage system should bind
- ◆ Specify the NIS domain name of your storage system
- ◆ Enable NIS on your storage system

You cannot configure the NIS slave during the setup procedure. To configure the NIS slave after the setup procedure is complete, you need to enable NIS slave by setting the option `nis.slave.enable` to On. For more information about enabling NIS slave, see “Enabling an NIS slave on your storage system” on page 107.

Data ONTAP interfaces to configure NIS

You can enable NIS and set NIS configuration values in either of these ways:

- ◆ Using FilerView
 - See “Configuring NIS with FilerView” on page 104.
 - You cannot use FilerView to configure the NIS slave.
- ◆ At the command line
 - See the appropriate instructions:
 - ❖ “Specifying NIS servers to bind to” on page 105
 - ❖ “Specifying the NIS domain name” on page 105
 - ❖ “Enabling or disabling NIS using the command-line interface” on page 105

If you want to use primarily NIS for host-name resolution, specify it ahead of other methods in the hosts map in the `/etc/nsswitch.conf` file. For information about editing the `/etc/nsswitch.conf` file, see “Changing the host name search order” on page 110.

Correct host-name resolution depends on the correct configuration of the NIS server. If you experience problems with host-name resolution or data availability, check the NIS server in addition to local networking.

For more information about your storage system’s NIS client, see the `na_nis(8)` man page.

Configuring NIS with FilerView

To set or modify NIS configuration values with FilerView, complete the following steps.

Step	Action	
1	In FilerView, click Network in the list on the left.	
2	In the list under Network, click Manage DNS and NIS Name Service.	
3	If you want to....	Then...
	Enable or disable NIS	Select Yes or No in the NIS Enabled field.
	Set or modify the NIS domain name	Enter a name in the NIS Domain Name field. Examples of configuration values are listed in “Specifying the NIS domain name” on page 105.
	Specify or modify NIS servers	Enter one or more IP addresses in the NIS Servers fields. Examples of configuration values are listed in “Specifying NIS servers to bind to” on page 105.

Enabling or disabling NIS using the command-line interface

To enable or disable NIS on your storage system, complete the following step.

Step	Action
1	Enter the following command: options nis.enable {on off} Use on to enable and off to disable NIS.

Specifying the NIS domain name

To specify the NIS domain name, complete the following step.

Step	Action
1	Enter the following command: options nis.domainname domain <i>domain</i> is the NIS domain name to which your storage system belongs; for example, typical NIS domain names might be <i>sales</i> or <i>marketing</i> . The NIS domain name is usually not the same as the DNS domain name.

Specifying NIS servers to bind to

You can specify an ordered list of NIS servers to which you want your storage system to bind. The list should begin with the closest NIS server (closest in network terms) and end with the furthest one.

To specify an ordered list of NIS servers you want your storage system to bind to, complete the following step.

Note

You can specify NIS servers by IP address or host name. If host names are used, make sure each host name, along with its IP address, is listed in the `/etc/hosts` file of your storage system. Otherwise, the binding with host name will fail.

Step	Action
1	<p>Enter the following command to specify the NIS servers and their order:</p> <pre>options nis.servers ip_address, server_name, *</pre> <p>The asterisk (*) specifies that broadcast is used to bind to NIS servers if the servers in the list are not responding. This is the default. If you do not specify broadcasting (that is, if you do not add the asterisk), and none of the listed servers is responding, NIS services are disrupted until one of the preferred servers responds.</p> <p>You can specify only IPv4 addresses or server names that resolve to IPv4 addresses using the /etc/hosts file on your storage system.</p> <p>Attention _____ Using the NIS broadcast feature can incur security risks.</p>

Example of specifying NIS servers to bind to: The following lists two servers and uses the broadcast default:

```
options nis.servers 172.15.16.1,nisserver-1, *
```

Your storage system first tries to bind to 172.15.16.1. If the binding fails, the storage system tries to bind to nisserver-1. If this binding also fails, the storage system binds to any server that responds to the broadcast. While bound to the NIS server that responded to the broadcast, the storage system continues to poll the preferred servers. As soon as one of the preferred servers is found, the storage system binds to the preferred server.

Enabling an NIS slave on your storage system

To enable an NIS slave on your storage system, complete the following step.

Step	Action
1	<p>Enter the following command:</p> <pre>options nis.slave.enable {on off}</pre> <p>Use on to enable the NIS slave and off to disable it.</p> <p>Note</p> <p>If the NIS slave is disabled, your storage system reverts back to the original configuration, in which it contacts an NIS server to resolve host names.</p>

Displaying NIS information

To display NIS information, complete the following step.

Step	Action
1	<p>Enter the following command:</p> <pre>nis info</pre>

For more information about the `nis info` command and resulting display, see the `na_nis(1)` man page.

You can display the following types of NIS information:

- ◆ NIS domain name
- ◆ Last time the local group cache was updated
- ◆ The following information about each NIS server that was polled by your storage system:
 - ❖ IP address of the NIS server
 - ❖ Type of NIS server
 - ❖ State of the NIS server
 - ❖ Whether your storage system is bound to the NIS server
 - ❖ Time of polling
- ◆ Information about the NIS netgroup cache
 - ❖ a. The status of the cache
 - ❖ b. The status of the "*" entry in the cache

- ❖ c. The status of the "*.nisdomain" entry in the cache
- ◆ Whether an NIS slave is enabled
- ◆ NIS master server
- ◆ Last time the NIS map was checked by the NIS slave
- ◆ NIS performance statistics:
 - ❖ Number of YP lookup network retransmissions
 - ❖ Total time spent in YP lookups
 - ❖ Number of network retransmissions
 - ❖ Minimum time spent in a YP lookup
 - ❖ Maximum time spent in a YP lookup
 - ❖ Average time spent in a YP lookup
- ◆ Response statistics for the three most recent YP lookups

Example:

The following example shows the statistics provided by the `nis info` command:

```
system1*> nis info
NIS domain is lab.netapp.com
```

```
NIS group cache has been disabled
```

IP Address	Type	State	Bound	Last Polled
Client calls	Became Active			

172.16.100.72	PREF	ALIVE	YES	Mon Jan 23 23:11:14 GMT 2006
0	Fri Jan 20	22:25:47	GMT 2006	

NIS Performance Statistics:

```
Number of YP Lookups: 153
Total time spent in YP Lookups: 684 ms, 656 us
Number of network re-transmissions: 0
Minimum time spent in a YP Lookup: 0 ms, 1 us
Maximum time spent in a YP Lookup: 469 ms, 991 us
Average time spent in YP Lookups: 4 ms, 474 us
```

3 Most Recent Lookups:

```
[0] Lookup time: 0 ms, 1 us    Number of network re-
transmissions: 0
[1] Lookup time: 5 ms, 993 us  Number of network re-
transmissions: 0
```

[2] Lookup time: 0 ms, 1 us Number of network re-transmissions: 0

NIS netgroup (*. * and *.nisdomain) cache status: Netgroup cache: uninitialized

*. * eCode: 0

*.nisdomain eCode: 0

NIS Slave disabled

NIS administrative commands

Data ONTAP supports the standard NIS administrative commands listed in the following table. For more information, see each command's man page.

Command	Function
ypcat	Prints an entire NIS map
ypgroup	Displays the NIS group cache entries
ypmatch	Looks up specific entries in an NIS map
ypwhich	Returns the name of the current NIS server

Changing the host name search order

How the host name search order is determined

If you use more than one method for host-name resolution, you must specify the order in which each name resolution service is used. This order is specified in the `/etc/nsswitch.conf` file in your storage system's root volume.

The default `/etc/nsswitch.conf` file

Data ONTAP creates a default `nsswitch.conf` file when you run the `setup` command on your storage system. The contents of the default file are as follows:

```
hosts: files nis dns
passwd: files nis ldap
netgroup: files nis ldap
group: files nis ldap
shadow: files nis
```

Note

Only the `hosts` entry in the `/etc/nsswitch.conf` file pertains to host-name resolution. For information about other entries, see the *Data ONTAP System Administration Guide* and the `na_nsswitch.conf(5)` man page.

By default, the host information is searched in the following order:

- ◆ `/etc/hosts` file
- ◆ NIS
- ◆ DNS

If you want to change this order, you can do so in either of these ways:

- ◆ By using FilerView
See "Changing the host name search order with FilerView" on page 111.
- ◆ By editing the `/etc/nsswitch.conf` file
See "Editing the `/etc/nsswitch.conf` file" on page 111.

Changing the host name search order with FilerView

To change the host name search order with FilerView, complete the following steps.

Step	Action
1	In FilerView, click Network in the list on the left.
2	In the list under Network, click Manage DNS and NIS Name Service.
3	In the Name Service section, select the desired values in the Hosts drop-down lists.

Editing the /etc/nsswitch.conf file

To change the order in which Data ONTAP searches for host information, complete the following steps.

Step	Action
1	If the /etc/nsswitch.conf file does not exist in your storage system's root volume, create it.
2	<p>Edit the file, entering each line in the following format:</p> <pre>map: service ...</pre> <p><i>map</i> for the host-name resolution service is hosts.</p> <p><i>service</i> is one or more of the following: files, dns, nis.</p> <p>For example, to change the resolution order to use NIS exclusively, change the hosts line to read as follows:</p> <pre>hosts: nis</pre>
3	Save the file.

Storage System Monitoring Using SNMP

5

About this chapter This chapter describes how Data ONTAP supports SNMP on your storage system and how you can use SNMP to monitor your storage system.

Topics in this chapter This chapter discusses the following topics:

- ◆ “Understanding SNMP implementation in Data ONTAP” on page 114
- ◆ “Managing the SNMP agent” on page 123
- ◆ “Creating SNMP traps” on page 129

Understanding SNMP implementation in Data ONTAP

SNMP process

If Simple Network Management Protocol (SNMP) is enabled in Data ONTAP, SNMP managers can query your storage system's SNMP agent for information (specified in your storage system's MIBs or the MIB-II specification). In response, the SNMP agent gathers information and forwards it to the SNMP managers using the SNMP protocol. The SNMP agent also generates trap notifications whenever specific events occur and sends these traps to the SNMP managers. The SNMP managers can then carry out actions based on information received in the trap notifications.

SNMP agent and MIB groups supported

For diagnostic and other network management services, Data ONTAP provides an SNMP agent compatible with SNMP version 1. This agent supports the MIB-II specification and the MIBs of your storage system. The following MIB-II groups are supported:

- ◆ System
- ◆ Interfaces
- ◆ Address translation
- ◆ IP
- ◆ ICMP
- ◆ TCP
- ◆ UDP
- ◆ SNMP

Note

Transmission and EGP MIB-II groups are not supported.

For more information about protocol support, see the `na_snmpd(8)` man page.

Types of traps in Data ONTAP

There are two types of traps in Data ONTAP:

- ◆ Built-in—Built-in traps are predefined in Data ONTAP and are automatically sent to the network management stations on the traphost list if an event occurs. These traps are based on one of the following:
 - ◆ RFC 1213, which defines traps such as coldStart, linkDown, linkUp, and authenticationFailure

- ◆ Specific traps defined in the custom MIB, such as diskFailedShutdown, cpuTooBusy, and volumeNearlyFull

For more information, see “Understanding traps in Data ONTAP” on page 116.

- ◆ User-defined—User-defined traps exist only after they are defined by a series of `snmp traps` commands or the FilerView SNMP Traps windows. These traps are sent using proxy trap ID numbers 11 through 18, which correspond to a trap’s MIB priority.

For more information, see “Creating SNMP traps” on page 129.

About the Data ONTAP MIBs

A Management Information Base (MIB) file is a textual description of SNMP objects and traps. Therefore, the Data ONTAP MIB files document the SNMP capabilities of the Data ONTAP version running on your storage system. MIBs are not configuration files—that is, values in the MIBs are not read by Data ONTAP, and changes to the MIB files do not affect SNMP functionality.

Data ONTAP provides two MIB files:

- ◆ A custom MIB (`/etc/mib/netapp.mib`)
See “Contents of the custom MIB” on page 119.
- ◆ An internet SCSI (iSCSI) MIB (`/etc/mib/iscsi.mib`)
See “Contents of the iSCSI MIB” on page 122.

Data ONTAP also provides a short cross-reference between object identifiers (OIDs) and object short names in the `/etc/mib/traps.dat` file. This is useful for creating user-defined traps, as discussed in “Defining or modifying a trap” on page 131.

Note

The latest versions of the Data ONTAP MIBs and `traps.dat` files are available online on the NetApp on the Web™ (NOW) site at <http://now.netapp.com/NOW/download/tools/mib/filer.shtml>. However, the versions of these files on the web site do not necessarily correspond to the SNMP capabilities of your Data ONTAP version. They are provided to help you evaluate SNMP features in the latest Data ONTAP release.

Understanding SNMP implementation in Data ONTAP

Understanding traps in Data ONTAP

About traps

Traps are mechanisms that alert you to significant events on your storage system. If SNMP is configured, traps are fired when a defined event, such as a network traffic interruption or line power failure, occurs. Trap information, in the form of MIB Object Identifiers (OIDs), is sent from your storage system's agent to an SNMP management station.

About built-in traps in Data ONTAP MIBs

Built-in traps in Data ONTAP MIBs are identified by the string TRAP-TYPE. For example, the following is a complete trap definition from the Data ONTAP custom MIB:

```
upsLinePowerOff          TRAP-TYPE
ENTERPRISE                netapp
DESCRIPTION
    "UPS: Input line power has failed and UPS is now on battery."

 ::= 142
```

Traps in the custom MIB are provided in a number of categories, including the following.

Category	Examples of trap messages
Disk Health Monitor	Degraded I/O, disk predictive-failure event
Disks	Disk - failure alert, shutdown, repaired
Fan	Fan - failed, shutdown, warning, repaired
Power supply	Power supply - failed, shutdown, warning, repaired
CPU	CPU busy, OK
NVRAM	Battery discharged, low
Cluster	Node failed, repaired
Volumes	Nearly full, full, repaired

Category	Examples of trap messages
Temperature	Over temperature, shutdown, repaired
Shelf	Fault, repaired
Global	Not recoverable, critical, not critical, OK
Soft quotas	Exceeded, normal
Autosupport	Send error, configuration error, successful send

Note

These categories are examples of the MIB trap contents; it is not an exhaustive list. The most complete listings are provided in the MIBs themselves.

About MIB trap priority

By convention, the right-most digit of a trap ID number indicates its priority (degree of severity), using the same enumeration as syslog entries. For example, trap ID 142 upsLinePowerOff is priority 2, alert.

Trap priorities are listed in the following table.

Trap ID last digit	Priority
1	emergency
2	alert
3	critical
4	error
5	warning
6	notification
7	information
8	debug

For more information, see the na_syslog.conf(5) man page.

**Where to get further
Information**

See the following RFCs for more information:

- ◆ RFC 1157—Defines and describes SNMP.
- ◆ RFC 1213—Defines and describes the SNMP MIB-II specification.
- ◆ RFC 1155—Defines and describes the structure and identification of management information for TCP/IP-based internets.
- ◆ RFC 1215—Defines the convention for defining traps for use with SNMP.
- ◆ RFC 1212—Gives concise MIB definitions.